HECC AWS Cloud: Overview

Starting summer 2019, we now offer the Amazon Web Services (AWS) Cloud on a pay-for-use basis. If you are a principal investigator (PI) directly affiliated with NASA, and you are interested in paying with NASA funding to use AWS Cloud for your projects, please contact us at support@nas.nasa.gov. Access to the AWS Cloud for users in a PI's project will be granted after funding is received and the AWS environment is configured for the PI's project.

Webinar: In addition to the articles in this section, you can find more information in our user training webinar, "Overview of HECC Pay-for-Use AWS Cloud." The recording and presentation slides are available in the HECC webinars archive.

This article provides a high-level overview of the HECC AWS Cloud. Follow the links in each section to learn more details about using the AWS Cloud.

HECC AWS Cloud Environment

As shown in this diagram, a project's HECC AWS Cloud environment includes both the HECC resources located at the NAS facility, and HECC's AWS resources. The resources at NAS are shared among projects, while those at AWS are private for each project. Some of the resources are optional, depending on the need of a project.

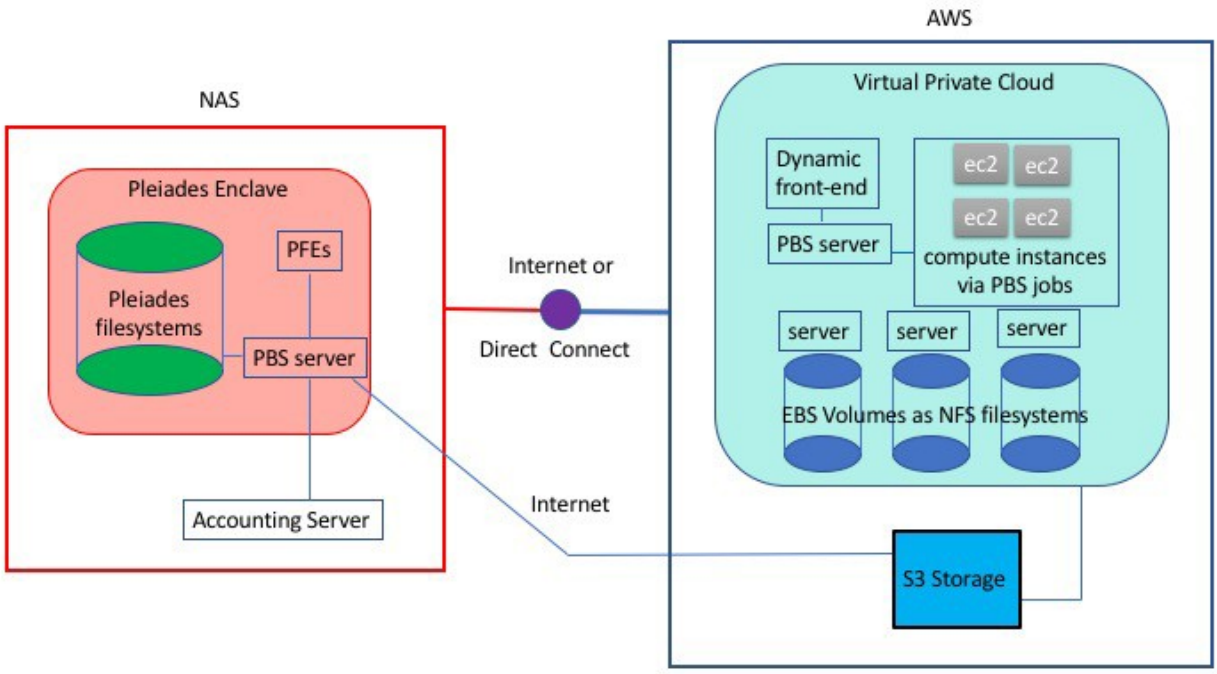

NAS-Located Resources

The purpose and functions of the NAS-located resources are described in this section.

Pleiades Front-End Systems (PFEs)

Each user in a project must have a NAS account in order to log in to a PFE. Access to AWS resources is accomplished via an SSH session from a PFE. Authentication to AWS is done behind the scenes.

From a PFE, you can remotely manage:

- PBS jobs submitted to AWS
- Data stored in AWS S3 storage

NAS PBS Server

A PBS server located at NAS, currently called clpbs-01, is used for accepting and checking PBS jobs submitted to run on AWS.

Pleiades /nobackup Filesystems

To run a job from NAS, you must submit it to AWS from your Pleiades /nobackup filesystem. The PBS output/error files of the AWS batch jobs are sent back to the $PBS_O_WORKDIR on your /nobackup filesystem.

Your /nobackup filesystem is also one of the sources or destinations for different types of file transfers between NAS and AWS.

Accounting Server

The accounting server manages your cloud allocations in Cloud Billing Units (CBUs). You can check your cloud allocation and usage by running the NAS **`acct_ytd`** and **`acct_query`** tools.

AWS Resources

The purpose and functions of the AWS resources are described in this section.

Note: The term *instance* used at AWS is logically equivalent to the term *node* used at NAS.

Dynamic Front End

In order to save costs, there is no static front end for AWS that runs 24x7. Instead, you launch a dynamic front end when you need one, and then shut it down when you're done. You only pay for the time that the dynamic front end is running.

A dynamic front end uses one AWS Elastic Compute Cloud (EC2) instance of your choice.

Use the dynamic front end to compile applications, and/or to manage and check PBS jobs or data on AWS. A limited number of software modules are available in the **`/nasa`** directory; if you need additional software, you must install it yourself under your own directory.

For more information, see AWS Elastic Compute Cloud (EC2) instance types.

AWS PBS Server

The PBS server at AWS coordinates with the PBS server at NAS to manage:

- Batch jobs submitted remotely from a PFE
- Batch jobs submitted locally from a dynamic front end

Note: Be sure to read AWS PBS Resources and Examples before submitting a batch job to AWS.

To save costs, the AWS PBS server is shut down when there are no batch jobs to manage.

The hostname of the AWS PBS server is different for each project. To find yours, use the following script on a PFE:

```
/u/scicon/tools/bin/aws_pbs_host
```

If you have multiple AWS projects, each with its own PBS server, add **`--group gid`** to the **`aws_pbs_host`** command to find the PBS server associated with that GID.

The instance type used for the server is currently m5.xlarge. It may change in the future.

Compute Instances for Batch Jobs

AWS offers many types of EC2 instances that you can use to run batch jobs. The default regions are **US West** for public cloud and **AWS GovCloud (US)** for government cloud. The HECC AWS Cloud uses the Linux AMI operating system.

Pricing varies with regions and operating systems.

For more information, see:

- EC2 instance types
- EC2 pricing
- EC2 regions and availability zones

Filesystem Servers

A filesystem requires a server with sufficiently high bandwidth to the filesystem. A c5.18x instance (similar to an Electra Skylake node) is typically chosen for the persistent filesystem described below. For a job-time filesystem, the instance chosen varies depending on the size to be allocated to the filesystem. For example:

- A c5d.18x instance for sizes below 1.8 terabyte (TB)
- An h1.16x instance for sizes above 4 TB and below 16 TB
- A c5.18x instance with Elastic Block Store (EBS) for all other sizes

To save costs, the server is shut down when the filesystem is not in use. If you need multiple filesystems, then you will need multiple servers.

Electric Block Store (EBS) Volumes as Filesystems

You can use Electric Block Store (EBS) Volumes as filesystems. AWS offers several EBS volume types and pricing. Depending on your project's need, different volumes can be provisioned into two types of filesystems:

- A persistent filesystem is independent of a batch job. Data in the filesystem persists even when the filesystem is not used. You pay for the persistent cost of the size provisioned, and for the time when the filesystem server is running.

  This type of filesystem is usually configured by HECC staff upon your request at the beginning of your project.
- A job-time filesystem is created at the beginning of a batch job and terminated at the end of the job. The data in this filesystem is lost at the end of the job. You pay for the size provisioned for the lifetime of the job, and the uptime of the filesystem server (if needed). Some job-time filesystems might not use EBS volumes.

For more information, see:

- EBS volumes
- EBS volume types
- EBS pricing

S3 (Simple Storage Service) for Long-Term Storage

Among the multiple S3 (Simple Storage Service) classes offered by AWS, HECC uses the standard class. Pricing depends on space used, the frequency of request, and the amount of data transferred out of S3 to the Internet. For more information, see:

- S3 storage classes
- S3 pricing

Article ID: 581
Last updated: 29 Apr, 2020
Revision: 44
Cloud Computing -> AWS Cloud -> HECC AWS Cloud: Overview
https://www.nas.nasa.gov/hecc/support/kb/entry/581/